

Steps Forward in the Philosophy of Information

Luciano Floridi

University of Hertfordshire

Department of Philosophy

l.floridi@herts.ac.uk

There is always much to learn when bright colleagues devote time and energy to evaluating one's own work carefully. The previous issue of *Etica & Politica / Ethics & Politics* offers a clear example of such an opportunity. The seven articles, constituting the Symposium on my book, *The Philosophy of Information*, cover a very wide spectrum of topics, insightfully and interestingly. They are very welcome steps forward in the development of the philosophy of information, not least because they make the latter much less a label for my own work and much more the name of an independent program of research, to which I have contributed. There is no space (and if there were, the exercise would be too tedious for the reader anyway) to detail all the lessons that I hope I might have learnt by reading them. So, what I shall do in this short article is to highlight some of the key points that I found most valuable. They are not the only ones, and other readers, or indeed the authors of the articles, may disagree on my selection; but they should offer a good sample of the stimulating contents to be found in the Symposium. I shall proceed by grouping the articles together not in terms of order of publication, but in terms of comments that I shall offer about them.

Let me start with the three more technical articles by Gamez, Cevolani, and Wolf. Gamez uses my work on semantic information to improve contemporary approaches to consciousness in terms of information integration. As Gamez acknowledges, and I agree with him, a major problem is that "consciousness is a real phenomena, whereas the information states that a system holds appear to largely depend on a subjective interpretation - and it is difficult to see how something that is metaphysically real can be correlated with or identified with a subjective interpretation". He seeks to resolve this impasse by arguing that the information integration theory of consciousness could be improved by revising it in terms of a *data* theory of consciousness, thus relying on my work on the notion of data in order to account for the objective properties of a physical system more successfully. This seems to me to be a very promising step forward, also in view of a further improvement in our understanding of self-consciousness.

Cevolani and Wolf also concentrate their attention on my work on semantic information, although not on how the latter accounts for data, but

on how it approaches the relation between data and information. Both concern my formulation of a theory of strongly semantic information, or TSSI (Floridi 2004; Floridi 2011a), according to which semantic information is well-formed, meaningful and truthful data. Each provides very useful improvements on my initial work. Cevolani concludes by suggesting that “the notion of partial truth captures all the essential intuitions underlying Floridi’s theory”. This may well be the case, and I remain indebted to Cevolani for his fine analysis of TSSI in terms of verisimilitude. Wolf offers an important “insight into some of the nuances of the theory [...]” and an improved “constraint that closely ties the informativeness function to the properties of the information system under study. This new constraint increases the robustness of TSSI”. In this case too, TSSI emerges as a better theory after Wolf’s careful revision. My contribution to their discussion of TSSI concerns a quick clarification of a philosophical feature of the theory that I am afraid I might have contributed to popularise, even if I do not endorse it. When presenting TSSI, I often speak in spatial terms, referring to distances, approximations, closeness, and so forth, between some information, say p , and its referent, that is, the state of affairs that the information in question seeks to capture, call it w . I do not believe this to be a problem in itself (see for example (Floridi 2010a)) but, in adopting such vocabulary, I might have given the impression that I subscribe to some kind of Platonism, whereby truths, truthful descriptions of w , or indeed features of w in themselves, are like targets, existing independently of our intellectual activities, fully formed, and hence such that some information p may miss them entirely, or get more or less closely, or perhaps hit. Such a Platonist reading would be a mistake, and I hope this explains why, in my research, I stayed away from any “verisimilitude” approach. From a constructionist perspective, truths are ultimately designed, so some semantic information p can be more or less close to its referent w only *metaphorically*, when discussing the foundations of our knowledge of the world in the deepest sense, or *secondarily*, when talking about the approximation between different bits of information among themselves, e.g. between my perception of where and when the train leaves and the train timetable. In everyday parlance, none of this generates any problem at all, but it is a philosophically troublesome issue for any philosopher who, like myself, rejects a representationalist interpretation of semantic information and therefore of the corresponding knowledge based on it. Ultimately, our perceptual interactions with reality consist in making sense of it, but such “semanticisation” is not closer to, or more distant from, the elements (data as constraining affordances) that it seeks to put together, in the same sense in which a *crème brûlée* is not more or less “distant” from its ingredients, if not just metaphorically. The analogy

clarifies that there is of course a perfectly reasonable sense in which our interpretations of the *interpreted* reality are more or less distant from their referents, again in the same sense in which this or that *crème brûlée* is more or less close to what we would consider a good *crème brûlée* according to the right recipe. But it is a mistake to assume that there is only one relation of reference: the rules of engagement between information and reality (dish vs. ingredients) are not the same as the rules of engagement between information and information (recipe vs. dish). In the former case, we are talking about the data-ingredients that come from (or constitute) reality and that information cannot approximate but only more or less correctly model. This is why, in *The Philosophy of Information*, I introduce a correctness theory of truth (Floridi 2010b; 2012). In the latter case, forms of correspondentism are perfectly justified, but then they are also philosophically naïve insofar as the fundamental game of making sense of the world is already over.

Wolf concludes his article by suggesting that “the analysis also provides a minor extension of TSSI into misinformation. Given the pervasiveness of misinformation and its interplay with information in the world in which we live, a real test of the strength of TSSI would be the development of a complete extension that thoroughly accounts for misinformation and its interplay with information. [...] As TSSI is refined to take into account experiences with information and misinformation, its value as an analytical tool will increase”. I agree and I would welcome such extension. It would perhaps address some of the criticisms moved by Fallis in his article. I must confess that I have the impression that Fallis is ultimately overstating his case. I do not think he is mistaken, but I still find perfectly fine to use “misinformation” to refer to false semantic information (e.g., Alice thinks that Milan is the capital of Italy, but she is misinformed), and “disinformation” to misinformation that is disseminated in order to mislead its receiver (e.g., Alice knows that the capital of Italy is Rome but she tells Bob that it is Milan in order to disinform him). The mis- in “misinformation” seems to go hand in hand with the mis- in “mistake”, but of course all this is largely a matter of linguistic conventions and uses. It is the point where philosophy becomes lexicography and taxonomy, and hence uninteresting. Of course, plenty of conceptual refinements are possible. After all, we should not expect the two concepts to be any less Protean than that of information itself (Floridi 2010c). But the important question to ask is not how far we can go with the fine slicing of a variety of different meanings and peculiar cases, but whether this is worth the effort. Fallis has not convinced me that it is. Imagine Alice measures a steel rail on a railroad and tells Bob that it is 20m long (66 ft in the UK). Bob double-checks her measurement and

discovers that, actually, to be precise, some steel rails can be quite a bit longer while others can be quite a bit shorter, and, moreover, their lengths can also vary depending on the temperature, and hence on when they are measured. So should Bob simply reject Alice's measure because "it is not yet clear that we have an adequate account" of the length of the steel rail? That depends on what the measure is for. For example, if Alice wishes to know how many steel rails her company needs to buy to cover the distance between two cities, centimetres will be rather pointless. However, they might be important when Bob deals with problems concerning the introduction of welded rail joints, which might require more finely grained measurements. The criticisms offered by Fallis in terms of visual disinformation, true disinformation, side effect disinformation, and evolutionary disinformation are interesting, but they concern centimetres, whereas I was giving the length of the steel rail in metres. Every concept should be used as precisely as possible, but not more precisely, to adapt something allegedly said by Einstein. Misinformation and disinformation are no exceptions. Nor is information, which remains a concept and a phenomenon impossible to pin down to those necessary and sufficient conditions so dear to some analytic philosophers who do not know better.

The previous comment leads me to Krebs' article and his welcome reminder that "Informativity [sic] depends on a given medial [sic] constellation as well as the interpreter, her competence and her contextually variable (epistemic) interests". This seems to be right, despite the unusual terminology. Concepts are semantic artefacts designed for a purpose, which they may fulfil more or less successfully. The point of conceptual analysis is not (and indeed should never be) to dissect them in a vacuum of contexts, practices, applications, and users, but to clarify, if necessary, how and why they work the way they do and, in case, try to improve them (Floridi 2008). There is something intellectually repulsive in displaying them in a glass-covered box, with the several conceptual species neatly labelled and mounted on pins.

I have left Beavers' and Durante's articles last because they seem to me to be the most encompassing. It would be pointless to go through the long list of points on which I agree with them. So let me offer in each case a small contribution to our mutual understanding.

Beavers rightly emphasises, among other things, that in my work "any philosophy of import – and let us not forget that those with no import are also not remembered – 1) is embodied in the ethos of a time and 2) set to answer questions that are pressing for its day. Consequently, 3) genuine philosophy changes over time". Today, this means developing a *constructionist* (not a *constructivist*) philosophy that can account for our

semantic artefacts and design or re-purpose those needed by our new infosphere. Let me add two clarifications to this. First, in the past, I have spoken of such a philosophical task in terms of *conceptual engineering* (Floridi 2011b). But more recently I came to realise that the very word “engineering” may generate confusion. True, both Descartes and Wittgenstein were engineers. Yet the concept has mechanical and deterministic overtones that I would be reluctant to endorse but that are difficult to shake off. I now find it rather clunky, and hence I much prefer speaking of *conceptual design*, especially in view of the fact that design is neither discovery nor invention, but indeed the art of exploiting constraining affordances intelligently and teleologically, in view of a specific goal. Thus, to the reader who may wish to build a network of references and “embed” the constructivism advocated in my philosophy of information into a more contextualised set of influences and perspectives, I would highly recommend reading (Deleuze and Guattari 1994). Second, once constructivism is understood as a realistic philosophy, which treats semantic artefacts as mind- *and* reality- dependent—in the same way as a house is the outcome of a specific architectural design both constrained and afforded by the building materials—it becomes clear that relativism is avoided at the expense of representationalism. For the equations in front of us are rather simple: we can either embrace a representationalist epistemology, which can avoid relativism by dropping the constructionist stance; or we can accept the fact that we are in charge of our conceptual constructions, some of which are very ill-conceived (astrology, homeopathy, Othello’s understanding of Desdemona’s behaviour, etc.) while others are increasingly successful in making sense of the world (astrophysics, medicine, the perfect understanding between Romeo and Juliet, etc.), but then constructionism without relativism becomes possible only by unveiling representationalist epistemology as another ill-conceived artefact.

In his article, Durante emphasises epistemic responsibility as a crucial factor in the constructionist stance supported by my philosophy of information. As I have just indicated, this is a very perceptive insight. It is one of the most important senses in which my philosophy of information is Cartesian and Kantian. Human beings have a special call, as the only semantic engines and intelligent informational organisms in the universe (at least to the best of current knowledge): they are the only entities responsible for the semanticisation of reality. This simple truth has many consequences. I would like to highlight two of them by way of conclusion of this short set of replies.

The first consequence concerns the nature of the constraining affordances. We saw that humanity designs semantic artefacts and

meaningful interpretations of reality by relying on constraining affordances. In terms of knowledge, these are the ultimate *data* that link our knowledge of reality to its referent. In terms of practices, such constraining affordances are the *rules* that shape our behaviour. Let me explain. In so far as data are relational differences (imagine the red light at a crossroads) and rules are conditional relations (“stop at the red light” understood as “if you do not stop at the red light you may cause an accident and incur in a fine”), they both enable agents like us (but not only like us, think of a robot, a dog, or a company) to interact with their environment and other agents more or less successfully. Data and rules are the constraining affordances thanks to which we are dynamically embedded in the world, to the extent that *anarchy* (the absence of any norms or rules) may be seen as the counterpart of *datalessness* (the absence of any distinctions or differences), leading to a kind of *agnosia* (the absence of any perception of data or rules because there are none).

The second consequence concerns the unique ethical duty that follows from humanity’s unique semanticising role. The epistemic responsibility involved in the design of a meaningful reality is not just an epistemological task, placed on our shoulders as individual epistemic agents. It is also, and probably more importantly, a social and ethical obligation that we have towards each other. If we and no one else make reality meaningful to ourselves and to others; if there is no other source of meaning in the universe but us; if our “semantic currency” is not backed up by some God standard; then there is only an immanent semantics, which is up to us to design, develop, protect, and share. This is our call. From it, it follows that each human life becomes valuable, and something to be cherished, as a precious source of sense-making. It would be a logical mistake to read such call solipsistically (and here is where I distance myself from Descartes and Kant), for the following reason. Semanticisation is an information process. Information, however, is not like any other ordinary good. It is *non-rivalrous*: Alice holding (consuming) the information that *p* does not prevent Bob from holding (consuming) the same information at the same time. It tends to be *non-excludable*. Some information – such as intellectual properties, non-public and sensitive data, or military secrets – is often protected, but this requires a positive effort precisely because, normally, exclusion is not a natural property of information, which tends to be easily disclosed and shareable. By contrast, if Alice lends her car to Bob, she cannot use it at the same time. Finally, once some information is available, the cost of its reproduction and dissemination tends to be negligible (*zero marginal cost*). This is of course not the case with many goods such as a loaf of bread. For all these reasons information may be sometimes seen as a *public good*, a view which in turn justifies the creation of public libraries or projects such as

Wikipedia, which are freely accessible to anyone. But then, semanticisation is a social process, to which we may contribute only a bit but from which we benefit enormously. Most, indeed *almost* (yet not) all the sense we can give to our lives is due to the sense-making activities of millions of other people. Hell is not the other, but the death of the other. As any old person knows, solitude is a social choice, made possible by the presence of others, but loneliness is a desperate condition due to the absence of any other.

To return to the initial comment on Durante's article, in the same way as data and rules are the relations representing the constraining affordances for our behaviour (not only epistemic), our semanticisation of them is both an epistemic and ethical task that we can fulfil only as social agents. Civilization is both an epistemic and an ethical concept for a multiagent system.

References

- Deleuze, G. and Guattari, F. (1994). *What Is Philosophy?* New York, Columbia University Press.
- Floridi, L. (2004). "Outline of a Theory of Strongly Semantic Information." *Minds and Machines* 14 (2): 197-222.
- Floridi, L. (2008). "Understanding Epistemic Relevance." *Erkenntnis* 69 (1): 69-92.
- Floridi, L. (2010a). "Information, Possible Worlds, and the Cooptation of Scepticism." *Synthese* 175: 63-88.
- Floridi, L. (2010b). "Semantic Information and the Correctness Theory of Truth." *Erkenntnis* 74 (2): 147-175.
- Floridi, L. (2010c). *Information - a Very Short Introduction*. Oxford, Oxford University Press.
- Floridi, L. (2011a). *The Philosophy of Information*. Oxford, Oxford University Press.
- Floridi, L. (2011b). "A Defence of Constructionism: Philosophy as Conceptual Engineering." *Metaphilosophy* 42 (3): 282-304.
- Floridi, L. (2012). "Semantic Information and the Network Theory of Account." *Synthese* 184 (3): 431-454.